# Out of the Cage of Shadows

Harald Baayen[a]
Shravan Vasishth[b],
Douglas Bates[c],
Reinhold Kliegl[b]

[a] University of Tübingen, Germany
[b] University of Potsdam, Germany
[c] University of Wisconsin-Madison, USA

arXiv:1511.03120v1 [stat.AP] 10 Nov 2015

Running Head: Out of the Cage of Shadows

Corresponding author:
R. Harald Baayen
Seminar für Sprachwissenschaft
Eberhard Karls University Tübingen
Wilhelmstrasse 19
Tübingen
e-mail: harald.baayen@uni-tuebingen.de

**Abstract**

Unlike molecules or plots of barley, subjects in psycholinguistic experiments are intelligent beings that depend for their survival on constant adaptation to their environment. This study presents three data sets documenting the presence of adaptive processes in psychological experiments. These adaptive processes leave a statistical footprint in the form of autocorrelations in the residual error associated with by-subject time series of trial-to-trial responses. Generalized additive mixed models (GAMMs) provide a unified framework within which both factorial predictors and covariates given with the experimental design, as well as non-linear random effects and interactions with experimental time can be uncovered and evaluated. GAMMs not only provide substantially improved fits to experimental data with time series structure, but also provide improved insight into predictors of theoretical interest, as well as a more refined window on the random effects structure. Our results challenge the standard advocated by Barr et al. (2013). The analytical cage of the maximal linear mixed model to which this standard confines the analyst is motivated by simulation studies which presuppose experimental data to be sterile, and free of any adaptive processes. However, when adaptive processes are present in real data, the simulation results of Barr et al. are no longer informative. For such data, the method of analysis cannot be purely design-driven, but must be in part driven by the data.

**Keywords**: generalized additive mixed models, factor smooths, within-experiment adaptation, autocorrelation, experimental time series

# 1 Introduction

Barr et al. (2013) proposed a new standard for the statistical analysis of experimental data. Based on a series of simulation studies, they argued that anti-conservative $p$-values are best avoided by fitting linear mixed effects models with all variance components included that could in principle be non-zero given the experimental design. Bates et al. (2015) and Matuschek et al. (2015b) pointed out that this standard may come with a loss of statistical power, and may result in overly complex models with parameters that are mathematically uninterpretable.

One might think that a potential loss of power and mathematically uninterpretable parameters are not too high a price to pay for an objective one-size-fits-all computational procedure that produces $p$-values that are assumed to be guaranteed not to be anti-conservative — an assumption the correctness of which, unfortunately, depends on the experimental design (Matuschek et al., 2015b). Proponents of this view would argue that the standard proposed by Barr et al. is useful for helping the field realize the importance of variance components other than just random intercepts, while at the same time avoiding problems that arise in the context of data-driven searches for optimal non-maximal models.

However, the simulations on which Barr et al. (2013) base their recommendation make very specific assumptions:

> We assumed no list effect; that is, the particular configuration of 'items' within a list did not have any unique effect over and above the item effects for that list. We also assumed no order effects, nor any effects of practice or fatigue. (p. 264)

In what follows, we will refer to data satisfying the assumptions of Barr et al. as 'sterile' data, data that are not infected in any way by hidden processes unfolding in time series of experimental trials.

In this study, we address the issue of whether the assumptions embodied in the simulations of Barr et al. can be justified for real data, as opposed to their simulated data, which are by design clean and sterile. If these assumptions can be shown to be unrealistic, at least for some real data sets, then it becomes questionable whether the standard proposed by Barr et al. has the proper empirical foundation to serve as a universal guideline for the application of mixed effect models.

In what follows, we discuss three data sets that are demonstrably infected and not sterile, and which illustrate why Barr et al.'s approach is not optimal. Indeed, we will show that when data sets are not sterile, a further problem arises for the methodological approach advocated by Barr et al., namely, that the purely design-driven approach to statistical analysis that they favor becomes infeasible. When data are not sterile, the analysis must, at least in part, be driven by the data.

To avoid misunderstandings, let it be clear that we have no objection whatsoever against factorial (or other) experimental designs. However, we take the position that factorial designs require true treatment factors. For instance, changing font size, luminosity, orientation of the stimulus, mood induction, etc., all are true treatments that are orthogonal to stimulus and subject properties. Conversely, situations in which a factorial design is suboptimal and comes with the risk of non-trivial loss of information comprise, e.g., research on lexical stimulus properties such as frequency of occurrence or acoustic duration. Although dichotomization into high versus low frequency, or long versus short acoustic duration, makes it possible to apply analysis of variance, dichotomization of continuous variables has many undesirable properties (Cohen, 1983; MacCallum et al., 2002; Baayen, 2010), and should be avoided if only to avoid unnecessary loss of power. For research questions addressing the consequences of such numeric stimulus properties for human processing, regression is the appropriate statistical tool (see also Balota et al., 2004).

Why might we expect that experimental data are not sterile? Because, unlike molecules or plots of barley, human subjects are complex organisms. Their survival depends on continuous adaptation to their environments, including the environments of psycholinguistic experiments.

In fact, it is well known that in time series of experimental trials, response variables such as reaction times elicited at time $t$ may be correlated with earlier reaction times at $t - k, k \geq 1$ (Broadbent, 1971; Welford, 1980; Sanders, 1998; Taylor and Lupker, 2001; Baayen and Milin, 2010). One source of temporal dependencies between trials is the presence of an autocorrelational process in the errors, potentially representing fluctuations in attention. Another source may be habituation to the experiment, possibly in interaction with decisions made at preceding trials (Masson and Kliegl, 2013). Alternatively, subjects may slow down in the course of an experiment due to fatigue. A further source of correlational structure in sequences of responses is learning. As shown by Marsolek (2008), the association strengths between visual features and object names are subject to continuous updating. Ramscar et al. (2010) and Arnon and Ramscar (2012) documented the consequences of within-experiment learning in the domain of language. Kleinschmidt and Jaeger (2015) report and model continuous updating in auditory processing in the context of speaker-listener adaptation. De Vaan et al. (2007) reported lexical decisions at trial $t$ to be co-determined by the lexicality decision and the reaction time to a prime that occurred previously at $t - 40$. Grammaticality judgements that change in the course of an experiment are reported by Dery and Pearson (2015). In what follows, we refer to the ensemble of learning, familiarization with the task, fatigue, and attentional fluctuations as adaptive processes.

When temporal autocorrelations are actually present in the data, but not brought into the statistical model, the residuals of this model will be autocorrelated in experimental time. The proper evaluation of model components by means of $t$ or $F$ tests presupposes that residual errors are identically and independently distributed. By bringing random intercepts and random slopes

into the model specification, clustering in the residuals by item or subject is avoided. However, such random slopes and random intercepts do not take care of potential trial-to-trial autocorrelative structure. Therefore the presence of autocorrelation in the residuals leads to imprecision in model evaluations and uncertainty about the validity of any significances reported. When strong autocorrelation characterizes the residuals, this uncertainty will render a model uninterpretable.

It might be argued that adaptive processes, if present, will have effects that are so minute that they are effectively undetectable. If so, the experimental design, and only the experimental design, would serve as a guide for determining the statistical model to be fitted to the data. Alternatively, one might acknowledge the presence of adaptive processes but claim that their presence gives rise to random and temporally uncorrelated noise. Any such adaptive processes would therefore be expected not to interact with predictors of theoretical interest. These two scenarios are fully compatible with the design-driven methodology advocated by Barr et al., provided of course that there is enough data to estimate all the parameters (Bates et al., 2015). If at all present, the effect of adaptive processes, being random, will be absorbed into the residual error, where it would be harmless.

However, it is conceivable that adaptive processes are present in a way that is actually not harmless. We distinguish two cases. First, adaptive processes may be present, without interacting with critical predictors of theoretical interest. In this case, measures for dealing with the autocorrelation in the errors will be required, without however affecting the interpretation of the experiment. Second, it is in principle possible that adaptive processes actually do interact with predictors of theoretical interest in non-trivial ways. If so, it is not only the autocorrelation in the residual error that needs to be addressed, but also the adaptive effects, which need to be disentangled from the main effects of the critical predictors.

In what follows, we discuss three examples of non-sterile data demonstrably infected by adaptive processes unfolding in the experimental time series constituted by the successive experimental trials. First, we re-analyze a data set with multiple subjects, and a $2 \times 2 \times 4$ factorial design with true treatments (Kliegl et al., 2015) and a single stimulus 'item'. We then consider a mega-study with auditory lexical decision (Ernestus and Cutler, 2015) using a regression design with crossed random effects of subject and item. Finally, we investigate a self-paced reading study in which subjects were reading Dutch poems, following up on earlier analyses presented in Baayen and Milin (2010).[1]

## 2 The KKL dataset

The experiment reported by Kliegl et al. (2015), a follow up to Kliegl et al. (2011), showed that validly cued targets on a monitor are detected faster than invalidly cued ones, i.e., spatial cueing effect (Posner, 1980) and that targets presented at the opposite end of a rectangle at which the cue had occurred were detected faster than targets presented at a different rectangle but with the same physical distance, an object-based effect (Egly et al., 1994). Different from earlier research, the two rectangles were not only presented in cardinal orientation (i.e., in horizontal or vertical orientation), but also diagonally (45 degrees left or 45 degrees right). This manipulation afforded a follow up of a hypothesis that attention can be shifted faster diagonally across the screen than vertically or horizontally across the screen (Kliegl et al., 2011; Zhou et al., 2006). Finally, data are from two groups of subjects, one group had to detect small targets and the other large targets. For an interpretation of fixed effects relating to the speed of visual attention shifts under these experimental conditions we refer to Kliegl et al. (2015).

[1] All analyses in this study were carried out with the help of the `mgcv` package (Wood, 2006, 2011) and the `itsadug` package (van Rij et al., 2015) for `R` (R Development Core Team, 2015).

Eighty-six subjects participated in this experiment. There were 800 trials requiring detection of a small or large rectangle and 40 catch trials. The experiment is based on a size (2) $\times$ cue-target relation (4) $\times$ orientation (2) design. Targets were small or large; rectangles were displayed either in cardinal or diagonal orientation, and cue-target relation was valid (70% of all trials) or invalid in three different ways (10% of trials in each the invalid conditions), corresponding to targets presented (a) on the same rectangle as the cue, but at the other end, (b) at the same physical distance as in (a), but on the other rectangle, or (c) at the other end of the other rectangle. Size of target was varied between subjects, the other two factors within subjects. The three contrasts for cue-target relation test differences in means between neighboring levels: spatial effect, object effect, and gravitation effect (Kliegl et al., 2011). Orientation and size factors are also included as numeric contrasts in such a way that the fixed effects estimate the difference between factor levels. With this specification the intercept estimates the grand mean of the 16 ($= 2 \times 4 \times 2$) experimental conditions. The data are available as KKL in the `RePsychLing` package at `https://github.com/dmbates/RePsychLing`. The dataframe already contains contrasts as numeric covariates. Dependent variables is the log of reaction time for correct trials completed within a 750-ms deadline. The total number of responses was 53765.

Bates et al. (2015) determined a parsimonious mixed model for these data, dealing with issues of overparameterization. We refitted this model using in addition a quadratic polynomial, which allowed us to include a well-supported nonlinear effect for stimulus onset asynchrony, which was varied randomly in an interval ranging from 300 to 500 ms in this experiment, but the effect of which had not been included in the initial LMM report of Bates et al. (2015).

Model criticism is an important but all too often neglected part of data analysis. Inspection of the residuals of the reference model reveals that although the residuals approximately follow a normal distribution, and although they are identically distributed, they are not independent.

The reaction times of a given subject constitute a time series, with experimental trial as unit of time. These trials can be ordered from the initial trial in the experimental list of trials, to the final trial in that list. In what follows, we refer to the time series of trials with the covariate `Trial` $= 1, 2, \ldots, k$. For the present experiment, we have 86 such time series, one for each of the 86 subjects. When we consider the residuals of the reference model, ordered by these time-series, we observe autocorrelative structure.

The strength of the autocorrelations in these by-subject time series varied from subject to subject, as illustrated for four exemplary subjects in the top panels of Figure 1. The autocorrelations for the subject in the top left panel are quite mild, and unlikely to adversely affect model statistics. For the second subject, we find evidence for autocorrelations up to at least lag 3. Autocorrelations increase for the third subject, and are still present at a lag of 15 trials. The subject in the rightmost panel show strong autocorrelations, indicating that a response time at trial $t$ is remarkably well correlated with the response at time $t - L$, for lags $L$ as large as 25.

Given that the residuals of the reference model (refitted with a GAMM) are not independent, it is unclear how reliable the estimates of model parameters and the assessments of the uncertainty about these estimates actually are. For this particular data set, strong autocorrelations such as for the last subject are exceptional, and hence it is likely that conclusions based on this model will be somewhat accurate. Nevertheless, a statistical model that is formally deficient is unsatisfactory, especially as there must be hidden temporal processes unfolding in this experiment that are not transparent to the analyst. Since the `KKL` data are clearly not sterile, a more fertile approach is to bring such hidden processes out in the open, and incorporate them into the statistical model.

Why are these autocorrelations present? In order to address this question, consider the plots in the lower set of panels of Figure 1. These panels present scatterplots of the data points for the four subjects in the corresponding top panels, to which two smoothers have been added, a loess

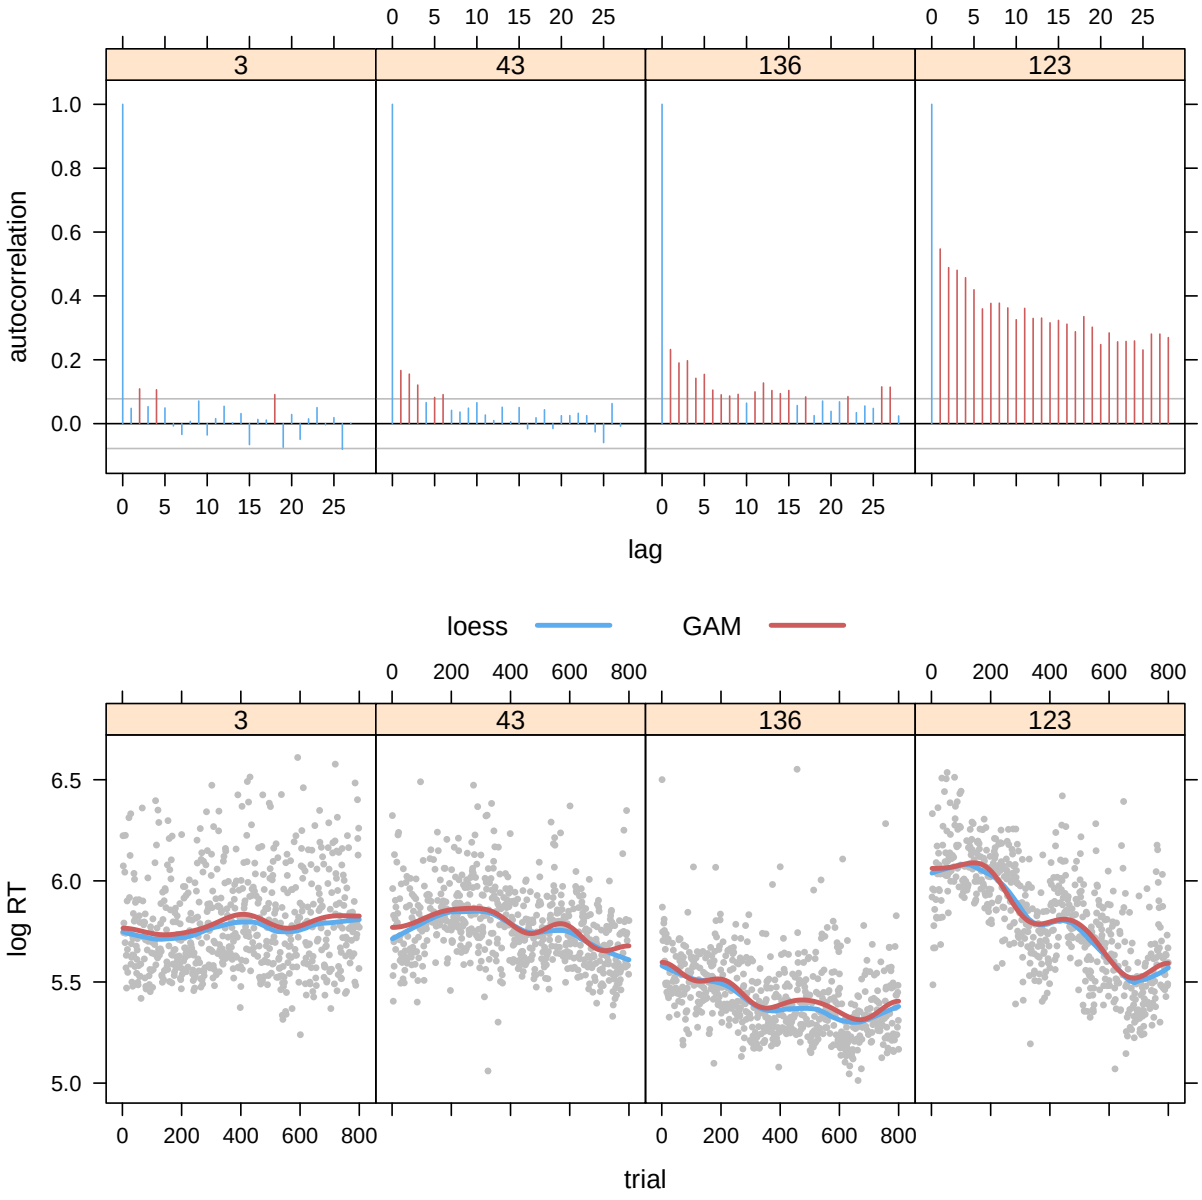


Figure 1: Autocorrelation functions for four subjects in the KKL dataset (upper panels), and the corresponding plots (lower panels) graphing log reaction time against trial, with a loess smoother (span = 0.2, in blue) and a GAM factor smooth (red).

smoother (blue) and a smoother obtained with a generalized additive model (red). For subject 3, whose time series of responses hardly shows any autocorrelation, we observe smooths that are close to horizontal lines. As the experiment proceeds, there are only very small changes in average response time. When we move further to the right in the array of panels, temporal patterns begin to emerge. As the experiment proceeded, subjects responded more quickly. Furthermore, it appears that there may be undulations in response speed. These oscillating changes in amplitude, if real, may reflect slow changes in subjects' attention or concentration over the course of the experiment.

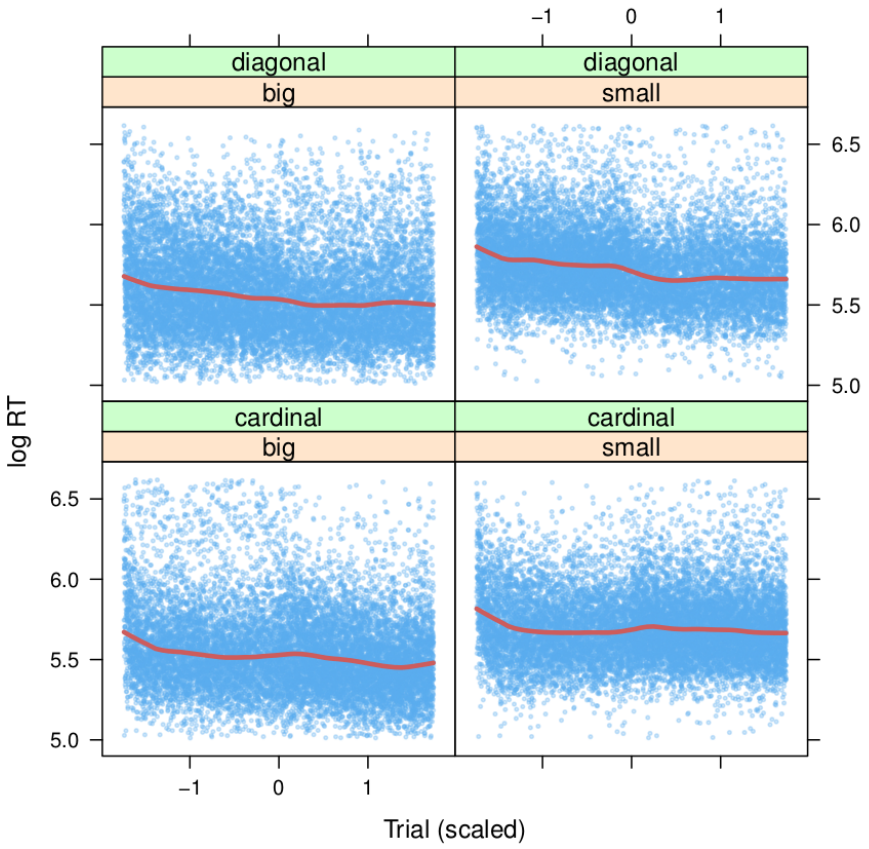


Figure 2: Loess smooths (span = 0.2) for the three-way interaction of `Trial` by `Orientation` by `Size`.

By contrast, the general downward trend present for subjects 43, 136, and especially 123 may point to familiarization with and gradual optimization of response behavior for the task.

The presence of a potential learning effect raises the question of whether learning proceeded in the same way across the different experimental conditions. Graphical exploration suggests that the rate at which subjects respond faster over time indeed varies, specifically so across the levels of `size` and `orientation`, as shown in Figure 2.

Within the context of the linear mixed model, the observed effects of trial can be taken into account by incorporating by-subject random slopes for `Trial`, and by allowing `Trial` to interact with `Size` and `Orientation`. As shown in Table 1, these extensions of our reference model are solidly supported by model comparisons using likelihood ratio tests. A summary of the final linear mixed model can be found in Table 3 in the appendix.

| | Df | AIC | logLik | deviance | Chisq | Df | Pr(>Chisq) |
|---|---|---|---|---|---|---|---|
| reference model | 25 | -25087.68 | 12568.84 | -25137.68 | | | |
| add Trial L * (sze+orn) | 28 | -25884.64 | 12970.32 | -25940.64 | 802.96 | 3 | $< 0.0001$ |
| add Trial Q * (sze+orn) | 31 | -26174.41 | 13118.20 | -26236.41 | 295.77 | 3 | $< 0.0001$ |
| add random slopes Trial | 33 | -26988.91 | 13527.45 | -27054.91 | 818.50 | 2 | $< 0.0001$ |

Table 1: Model comparisons for linear mixed models fitted to the `KKL` dataset.

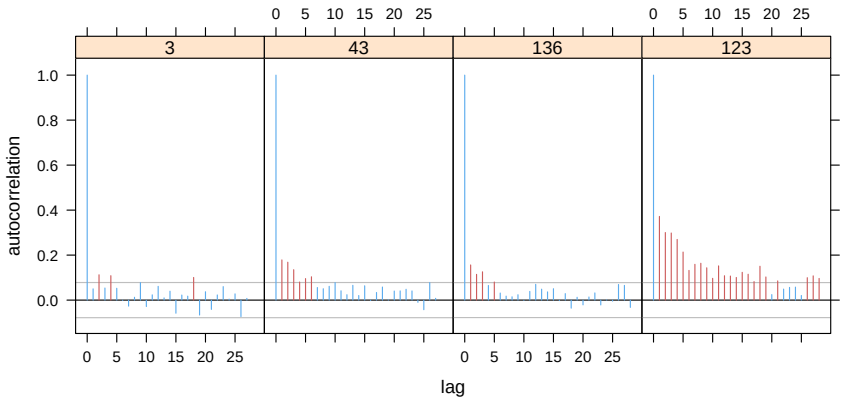


Figure 3: Autocorrelations in the residuals of the extended linear mixed model fitted to the `KKL` data.

Figure 3 presents the autocorrelation functions for the residuals of this final, comprehensive, linear mixed model. Comparison with the top panels of Figure 1 shows that for subjects 136 and 123, the autocorrelation in the residuals has been reduced substantially, thanks to bringing the effects of `Trial` into the model. Nevertheless, some autocorrelation remains present.

For further reduction of autocorrelations, it is necessary to relax the assumption that the subject-specific effects of `Trial`, currently modeled by means of by-subject random intercepts and random slopes, are strictly linear. The smooths presented in the bottom panels of Figure 1 suggest that undulations may ride on top of the linear trends. What we need, then, is a way of relaxing the linearity assumption for the by-subject random effects of `Trial`. The generalized additive mixed model (GAMM, see, e.g., Hastie and Tibshirani, 1990; Lin and Zhang, 1999; Wood, 2006, 2011, 2013) provides such a nonlinear counterpart to the combination of random slopes and random intercepts, in the form of penalized factor smooths. A factor smooth for `Trial` by subject sets up a separate smooth for each level of the factor `Subject`. When we add the constraint that each smooth should have the same smoothing parameter, and penalize the smooths for wiggliness, thereby shrinking them towards zero, we obtain 'wiggly random effects'.

The red smooths in Figure 1 are such factor smooths. They are very similar to the loess smooths, but are slightly more sensitive to the undulations in the data. Althought it might seem there is a risk that the factor smooths are modeling noise rather than signal, this is unlikely as the factor smooths are evaluated within the general framework of the generalized additive mixed model, and hence it is possible to assess whether they contribute significantly to the model fit. By way of example, across 100 random permutations of `Trial` for the four subjects of Figure 1, a significant factor smooth was obtained in 3 instances for $\alpha = 0.05$ and for zero cases for $\alpha = 0.01$, indicative of nominal Type I error rates. This example illustrates informally that factor smooths are unlikely to find and impose nonlinear structure when there is none.

Importantly, we think the undulating random effects captured by factor smooths represent the ebb and flow of attention. They emerge not only in the present data set, but have been observed for visual lexical decision (Mulder et al., 2014) as well as for word naming and for EEG data (Baayen et al., 2015). If this interpretation is correct, penalized factor smooths are the appropriate statistical tool to use. We explicitly do not want to model these fluctuations in attention as fixed effects, because there is no reason to believe that if the experiment were replicated, a given subject would show exactly the same pattern. It is more realistic to expect that changes in attention will again

be present, with roughly the same magnitude, but with ups and downs occurring at different points in time. In other words, we are dealing here with temporally structured noise, and the penalized factor smooths make it possible to bring such 'wiggly random effects' into the statistical model.

| | AIC | fREML | edf | $p$ |
|---|---|---|---|---|
| reference model | -26012.06 | -12500.67 | 24.00 | $< 0.0001$ |
| linear model | -28047.29 | -13422.47 | 34.00 | $< 0.0001$ |
| factor smooths | -30774.94 | -14449.56 | 25.00 | |
| smooth Trial | -31040.29 | -14582.62 | 33.00 | $< 0.0001$ |

Table 2: Model comparison for GAMMs fitted to the `KKL` dataset. (The reference and linear models were refitted with GAMMs to ensure comparability across all four models.)

Table 2 lists four GAMMs that we fitted to the `KKL` data. The first is the reference model, but refitted with GAM software (the `mgcv` package for `R`), rather than with LMM software (the `lme4` package for `R`), with as the only change that a smoothing spline is used for the `SOA` covariate, instead of a quadratic polynomial. The second model has the same specification as the final linear mixed model (summarized in Table 3 in the appendix), but refitted with a GAMM. (These two models were refitted because both the estimation algorithms and the way in which degrees of freedom are handled differ between `lme4` and `mgcv`.) As expected, the linear model, which includes effects for `Trial`, outperforms the reference model. The third model, which replaces the by-subject random intercepts and slopes by factor smooths, provides a better fit with fewer effective degrees of freedom. Addition of the three-way interaction of `Trial` by `Size` and `Orientation` improves the model further.

Figure 4 visualizes this three-way interaction of `Trial` by `Orientation` by `Size`. The effect of learning is larger for the cardinal presentation (upper panels) than for the diagonal presentation (bottom panels). For both large (left panels) and small (right panels) stimuli, we observe rapid initial learning, which levels off more for small than for large stimuli. Big stimuli with diagonal presentation elicited the most gradual accommodation pattern, with response times gradually becoming shorter.

Figure 5 clarifies that the full GAMM succeeded in further reducing the autocorrelations in the residuals. This reduction is due almost exclusively to the use of factor smooths, with only tiny amelioration by adding in the three-way interaction with `Trial`. This result is important for two reasons. First, it is unlikely that the removal of autocorrelation in the residuals could be accomplished by factor smooths fitting noise rather than signal. Likewise, it is unlikely that the huge reduction in AIC when going from the linear mixed model to the GAMM (2728, see Table 2) could be accomplished by just fitting noise. Second, the presence of slowly undulating processes in experimental data is a phenomenon that is itself of theoretical interest, and invites interpretation, clarification, and replication.

The remaining autocorrelations that are visible in Figure 5 are unlikely to be harmful, but to play safe, one might consider removing subject 123 from the dataset and refitting the model. Alternatively, these last remaining autocorrelations might be due to a simple AR(1) autocorrelative process in the errors, according to which the current error is linearly proportional to the preceding error with a proportionality constant $\rho$ plus Gaussian noise. Pinheiro and Bates (2000) and Gałecki and Burzykowski (2013) provide extensive discussion of how autocorrelation processes can be accounted for within the mixed modeling framework. With a mild proportionality constant $\rho = 0.15$, autocorrelations are almost completely removed. Below, we will discuss the use of this parameter in more detail.

Addressing the autocorrelation issue for the `KKL` data set does not lead to major changes in significances and magnitudes of fixed-effect coefficients and the magnitudes of the (significant) coefficients, as reported in Bates et al. (2015) and Kliegl et al. (2015). Nevertheless, the GAMM

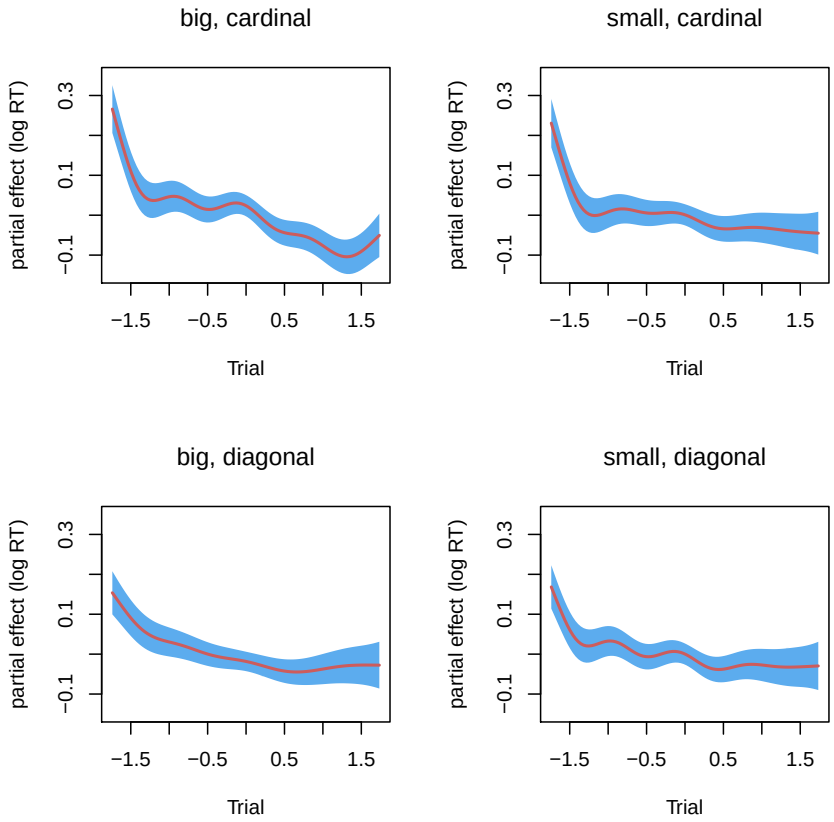


Figure 4: The `Trial` by `Orientation` by `Size` interaction in the full GAMM model for the `KKL` data.

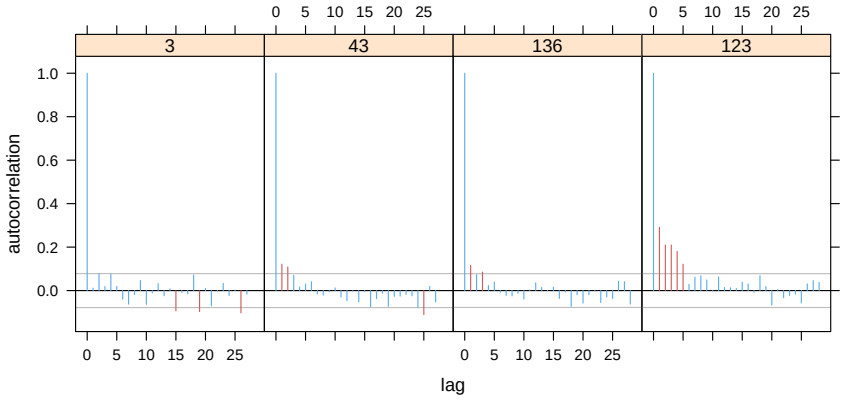


Figure 5: Autocorrelation functions for four subjects in the full GAMM fitted to the `KKL` data.

offers enhanced insight into the data, specifically with respect to the effect of `Orientation`. In the reference model, the coefficient of this main effect was estimated at 0.041, with a standard error of 0.010 ($t = 3.9$). However, in the full linear mixed effect model, the coefficient is smaller (0.014), comes with greater uncertainty (standard error 0.09), and is unlikely to be significant ($t = 1.5$). However, the final GAMM estimates the coefficient at 0.039, with a standard error of 0.016 and a $t$

value of 2.491, reporting $p = 0.013$. Given the interaction of `Trial` by `Orientation`, the greater uncertainty about this main effect makes perfect sense. Furthermore, as expected, the variance of the BLUPs for the by-subject random effects of `Orientation` in the reference model is larger (by a factor two) than in the final GAMM ($p < 0.0001$, $F$-test). Again, this makes perfect sense, as part of what originally looked like random noise linked to orientation can now be attributed to a learning effect over experimental time.

In summary, the `KKL` dataset violates the assumption underlying the proposed standard of Barr et al., in that this experimental data set is not sterile, but infected by the 'human factor'. The by-subject time series are characterized by autocorrelated errors. Unlike particles in physics, or plots of barley in agricultural experiments, human subjects are intelligent beings whose behavior is not random over time, but adaptive. The present reanalysis shows that subjects adapt in different ways to the novel manipulation of canonical versus diagonal positioning of visual stimuli, which is a theoretically fertile result. The proposed standard of Barr et al., however well-intended, has the unfortunate side-effect of locking analytical practice in a methodological cage from which the true structure experimental data, rich and fertile in perhaps unexpected ways, cannot be appreciated.

# 3 The baldey dataset

Our second example addresses the analysis of the response latencies elicited in the auditory lexical decision megastudy of Ernestus and Cutler (2015). Ernestus and Cutler include in their statistical analysis the reaction time to the preceding trial, as a way of controlling for temporal dependencies in by-subject time series. As shown by Baayen and Milin (2010), the inclusion of preceding reaction time successfully removes a considerable amount of autocorrelation in the residuals. Ernestus and Cutler also included `Trial` as a main effect, together with by-subject random slopes for `Trial`.

In what follows, we present an analysis of the reaction times of the `baldey` data, which shows that the human factor is even stronger in these data than suggested by the analyses of Ernestus and Cutler.

In our analysis — which is far removed from a "comprehensive" analysis of this rich data set — we departed from the analyses of Ernestus and Cutler in several ways. First, we analyzed an inverse transform of the reaction times (-1000/RT) rather than a logarithmic transform, as both graphical inspection and an analysis following Box and Cox (1964) indicated the inverse transformation to better approximate normality.

Second, including the preceding reaction time is a suboptimal way of bringing autocorrelation in the errors under control. We therefore explored by-subject factor smooths in combination with the possibility of an AR(1) process in the errors.

Third, we relaxed the assumption that effects would be the same across the two sexes. There are indications that males and females may be differentially sensitive to word frequency (Ullman et al., 2002; Balling and Baayen, 2008), but a sex by frequency interaction is not always found (Balling and Baayen, 2012; Tabak et al., 2005, 2010). As the `baldey` data set combines a perfectly balanced set of subjects (10 males and 10 females) with a large number of items (2780 Dutch words), it provides an excellent testing ground for differential effects of the two sexes in lexical processing.

Finally, we relaxed linearity assumptions, replacing a strictly linear mixed model by a generalized additive mixed model.

The random effects structure of the model for the reaction times for words (see Table 5 in the appendix for a statistical summary) included random intercepts for word, as well as by-word random intercepts for gender. Different words enjoy different popularity across the sexes (see also Baayen et al., 2015), and adjusting by-word intercepts for gender results in a tighter model fit. With

respect to subject, we included by-subject factor smooths for session.[2] The data for this mega-study were collected over 10 sessions, and once by-subject factor smooths for session were included in the model, by-subject factor smooths for `Trial` became redundant. For subject, random slopes for the acoustic duration of the stimulus word were also well supported.

Lemma frequency (the summed frequency of a word's inflectional variants) revealed a non-linear effect that differed between females and males, as shown in the top panels of Figure 6. Furthermore, the effect of the acoustic duration of the auditory stimulus was different for females and males, and showed small but statistically well-supported modulation by `Trial`. This interaction is visualized in the bottom panels of Figure 6. About two-thirds through an experimental session, the effect of acoustic duration decreased, for both males and females. (In Figure 6, this can be seen by noticing the reduced gradient for (scaled) `Trial` equal to 0.5: the number of contour lines crossed when moving horizontally across the plot, i.e., for increasing acoustic duration, is smaller than early on in the experiment.)

As for the `KKL` dataset, inclusion of session and trial in the model did not absorb all autocorrelation in the residuals. With $\rho = 0.2$, the remaining autocorrelations were properly accounted for.

What this analysis shows is that subjects participating in an experiment with language materials bring with them their own — in the present case, gender-specific — experiences with the language, and that these specific experiences will lead to differentiated effects, not only of well-established predictors such as word frequency and acoustic duration, but also for effects of learning and fatigue. Again, the human subject emerges as a moving target (see also Ramscar et al., 2013).

# 4 The poems dataset

Our final example addresses a data set previously discussed by Baayen and Milin (2010), available under the name `poems` in the `RePsychLing` package. This data set comprises a total of 275996 data points from 326 subjects, for 2315 words appearing across 87 modern Dutch poems. Words are partially nested under poems. Any given subject read only a subset of poems.

Baayen and Milin included random intercepts for subject, word, and poem, as well as several by-subject and by-word random slopes for various numerical predictors. These authors sought to eliminate the problem of autocorrelated errors by including trial as a predictor, as well as the self-paced reading latency at the preceding word.

As discussed in detail by Bates et al. (2015), the model of Baayen and Milin is overparameterized with respect to its random effects structure. Given the Zipfian shape of word frequency distributions and the large numbers of words occurring only once in the corpus of poems, data are too sparse to include word as random-effect factor. Furthermore, correlation parameters for by-word random intercepts and slopes in the Baayen and Milin model were quite large, with absolute magnitudes $> 0.8$, often an indicator of an overparameterized model. Furthermore, as for the `baldey` data discussed in the preceding section, including an AR(1) process in the errors is a more principled and effective solution than including the previous reaction time as covariate.

Within the context of the present discussion, the `poems` dataset is of interest for two reasons. First, because subjects are reading connected discourse rather than responding to unrelated isolated stimuli, the autocorrelation in their responses is much stronger than in the `KKL` and `baldey` datasets. This is illustrated in the top panels of Figure 7 for four exemplary subjects. In this dataset, there

[2] Given the small number of sessions (11) and the large number of observations for each session (around 4500 for each subject), one could opt for treating session as a factor, and including by-subject random slopes for this factor. Results are very similar, with the factor smooths showing slightly more shrinkage.

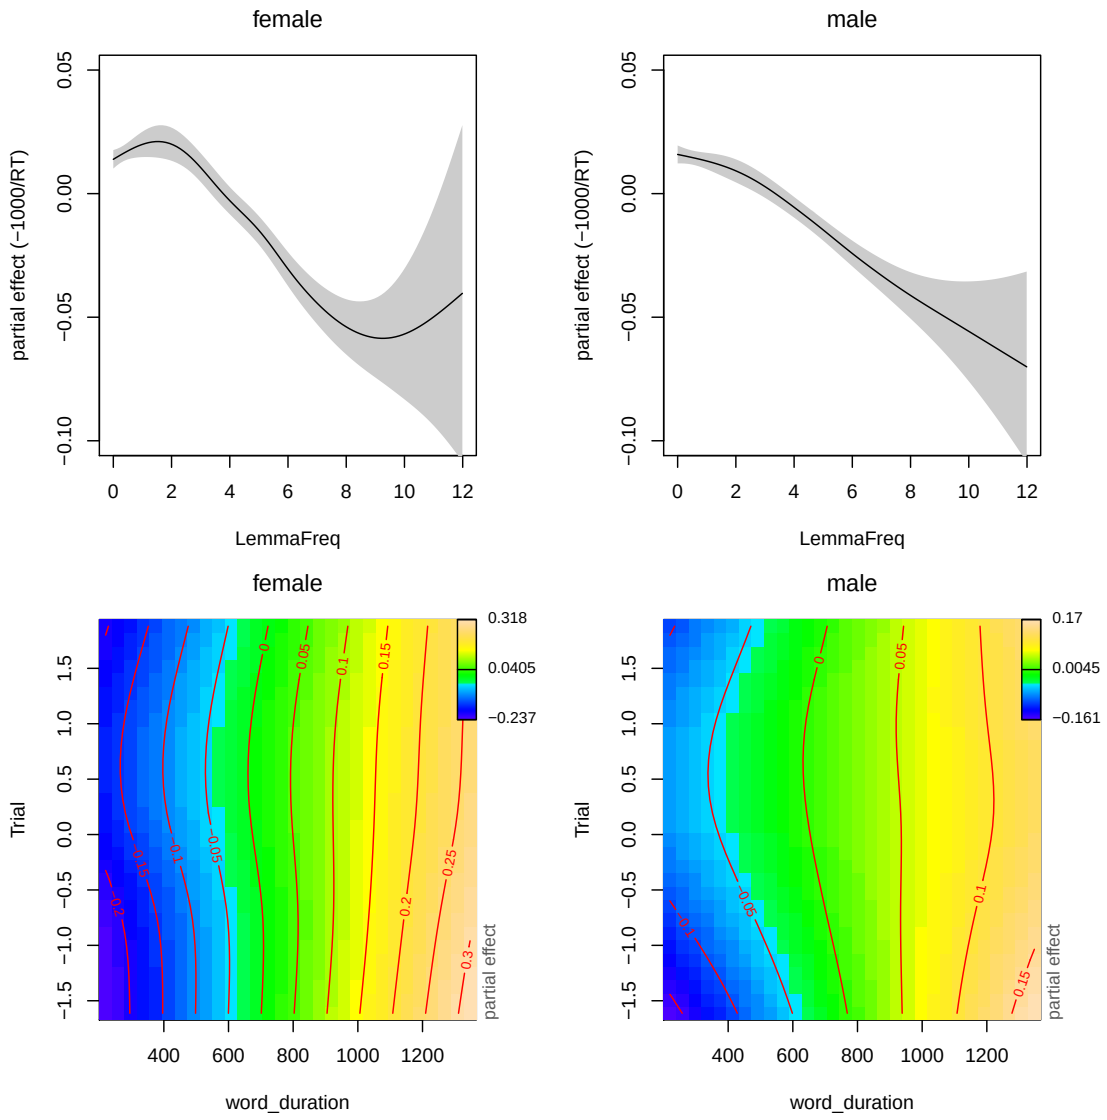


Figure 6: Effects of `Trial`, `Frequency`, `Word Duration`, and `Sex` in the `baldey` data set.

is only a handful of subjects without autocorrelations, and there are subjects with even stronger autocorrelations than the ones shown here. The second row of panels shows the corresponding scatterplots with loess and GAM smoothers. Especially for subjects 19 and 183, there are temporally concentrated spikes of long reading times that are beyond what a GAM smooth can capture. The lower set of panels illustrates the limitations of what the GAMM fitted to this data set (and described in detail in Table 6 in the appendix) can accomplish. For subject 265, the autocorrelations are properly removed, and for subject 176, the reduction in autocorrelation is perhaps satisfactory. This is not the case for subjects 19 and 183, unsurprisingly given the spiky trends in the scatterplots.

Increasing $\rho$ is not an option. As discussed in further detail in Baayen et al. (2015), since different subjects typically emerge with different degrees of autocorrelation, one would want to adjust the $\rho$ parameter for each individual subject. Unfortunately, it is at present not known how to achieve this mathematically. As a consequence, the optimal $\rho$ is one that strikes a balance, such that the autocorrelation for subjects with strong autocorrelation is reduced as much as possible,

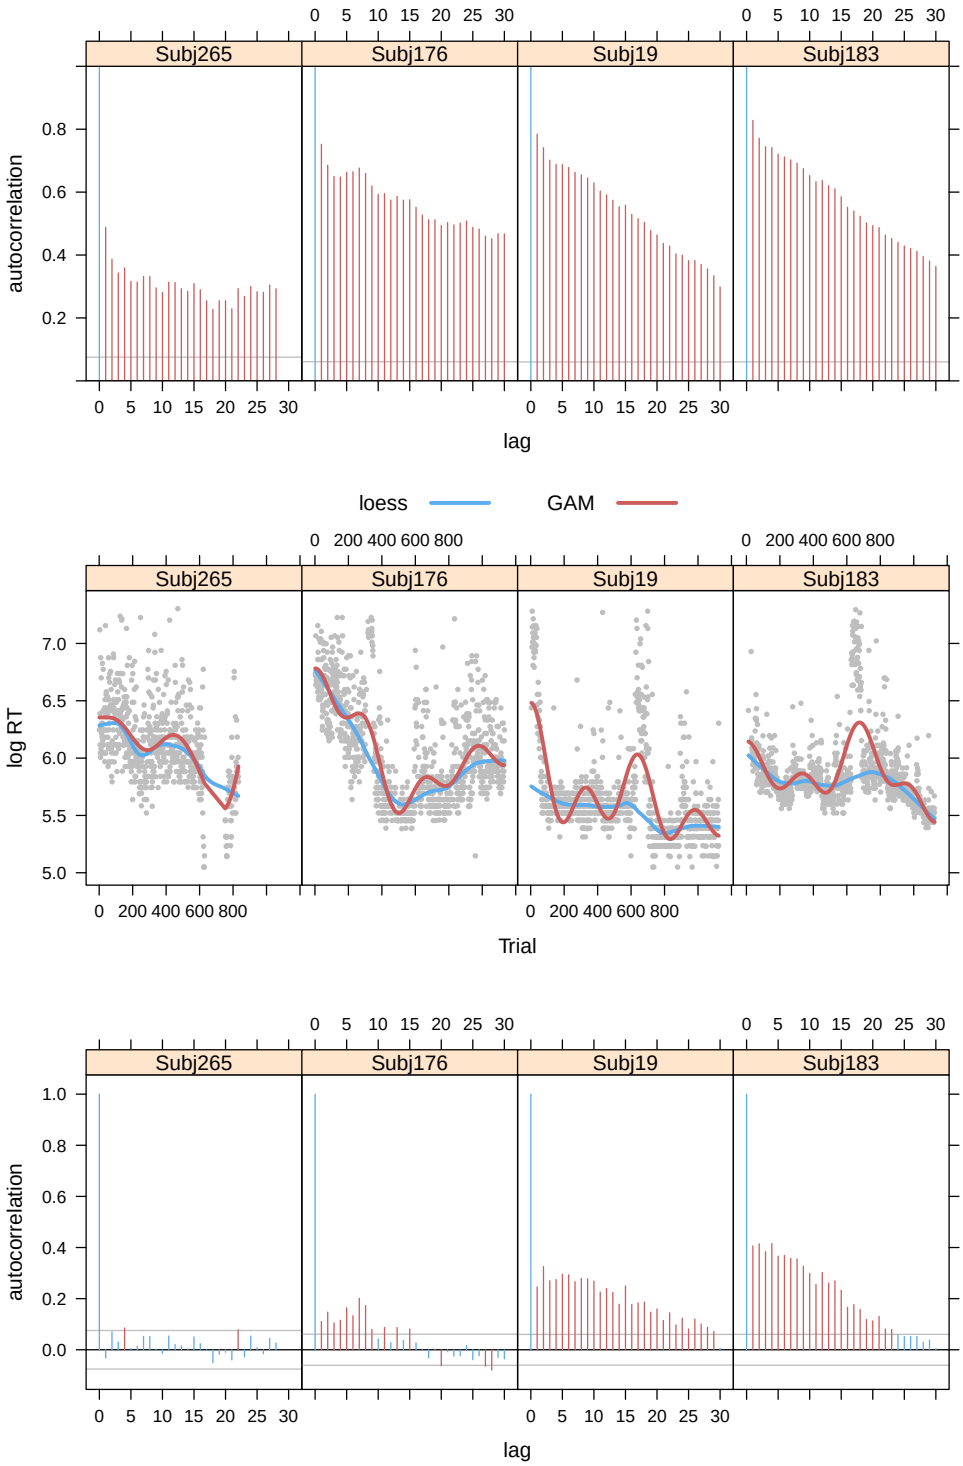


Figure 7: Autocorrelation functions for log reaction time for four exemplary subjects in the `poems` data set (upper panel), the corresponding plots (center panels) graphing log reaction time against trial, with a loess smoother (span = 0.2, in blue) and a GAM factor smooth (red), and autocorrelations in the residuals of the GAMM ($\rho = 0.3$).

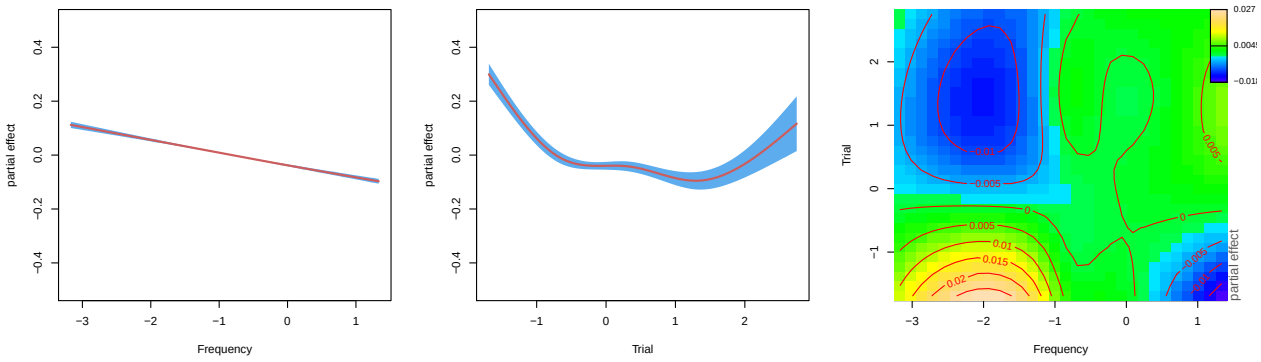


Figure 8: Partial effects for `Frequency` and `Trial` in a GAMM with an ANOVA decomposition into additive main effects and interaction, for the `poems` data set.

without introducing artificial negative autocorrelation at short lags for subjects with little or no actual autocorrelation in their residuals.

Keeping in mind the caveat that the GAMM provides an imperfect window on the complex quantitative structure of the `poems` data, it is of interest that word frequency appears to enter into a strong interaction with `Trial`. The appendix reports two models, one with a single multivariate smooth for these two predictors, and one in which their joint effect is decomposed into separate, additive, main effects of `Trial`, `Frequency`, and their interaction (see also Matuschek et al., 2015a). These three partial effects are presented in Figure 8. We see a linear facilitatory main effect of (log-transformed) `Frequency` (left panel), a U-shaped effect of `Trial` (center panel), and an interaction that rides on top of these two main effects (right panel). The contour plot indicates that in the early trials, frequency had a more downward-sloping gradient. Later in the experiment, the effect of frequency is attenuated. The reduction in the magnitude of the frequency effect as the experiment proceeds makes sense. As subjects read through the poems selected for them, they tune in to this particular genre and its vocabulary. Words are repeated, and become more predictable as words align into sentences, and sentences into poems. As a consequence, frequency of occurrence as a contextless lexical prior becomes increasingly less informative.

## 5 Concluding remarks

The Open Science Collaboration (2015) reported the results of an extensive series of replication studies, and documented a 50% drop in observed effect sizes and a drop from 97% to 36% of significant results. Publication bias (Francis, 2012), non-random selection of stimuli and subjects (Forster, 2000; Sander and Sanders, 2006; Francis, 2013), societal changes (Ramscar et al., 2015) and especially lack of power (Button et al., 2013; Westfall et al., 2014) are major concerns. Clearly, many of the results in the published literature are but moving shadows of the true effects, which raises the question how to escape from the current methodological cave.

The paper by Barr et al. proposes the imposition of a one-size-fits-all standard that is supposed to protect the field against anti-conservative p-values. The present study documents three data sets for which this proposed standard is unhelpful and even misleading, and has the effect of converting the cave of shadows into a cage of shadows. Unlike molecules or barley, the units from which response variables in psychology are harvested are intelligent beings that constantly adapt to their

environment. Humans learn. When adaptation is demonstrably present in experimental data, researchers will want to bring associated processes under statistical control. Failure to take the human factor into consideration comes with the risk of misunderstanding the quantitative structure of the data and the extent to which this structure is shaped by the predictors of interest.

We have illustrated how generalized additive mixed models can be helpful for obtaining a more nuanced understanding of experimental data. We do not claim that with GAMMs researchers will finally emerge from shadows and apprehend the true effects themselves. But we are finding GAMMs helpful in sharpening the contours and outlines of these effects. For almost all data sets that we have investigated with GAMMs, we have obtained better fits by including by-subject factor smooths for trial. For a subset of our data, further interactions of trial with predictors were present, of which examples were provided above.

One important methodological consequence of taking the human factor seriously is that statistical analysis, unavoidably, becomes partly data driven. Given our current state of knowledge, it is often impossible to formulate concrete hypotheses about how adaptation, learning, attentional fluctuation, and effects of fatigue, might interact with predictors of interest. All we can do is allow the data to speak for themselves, by bringing experimental time as covariate into the model and by checking for — potentially nonlinear — interactions with the predictors of theoretical interest.

We do not claim that generalized additive mixed models provide the solution to all problems. As any statistical model, GAMMs build on assumptions that are hoped to be reasonable, but of which we often know they involve substantial simplifications. The penalized factor smooths that we have used, for instance, assume a common smoothing parameter for all subjects, which may be true, but may also be incorrect. The factor smooths build on spline theory, but when time series of reaction times are spiky instead of smoothly wiggly, they are better than nothing, but certainly not perfect. Furthermore, the adjusting of the errors for an autocorrelative process can be much too simplistic in two ways. First, as we have demonstrated, $\rho$ should ideally vary with participant, which is currently not possible. Second, there is no guarantee that the autocorrelative process is a simple AR(1) process — dependencies may well stretch further back in time. Nevertheless, we believe that with GAMMs, researchers have a tool in hand — fortunately a tool that is under rapid development — with which the real-life complexities from which psycholinguistic data suffers can start to be investigated.

One of the considerations in favor of standards such as proposed by Barr et al. is that all subjectivity is taken out of the analytical procedure. In the approach advocated here, however, responsibility is transferred from the experimental design back to the researcher. Rather than deploring this, we believe the way forward is to publish one's data. Whereas before the advent of the internet, reporting $p$-values in printed journals was the only way to make an argument that a particular effect is present, current communication technology makes it possible to publish not only $p$-values but the data themselves, using platforms such as, e.g., the Mind Research Repository at `http://openscience.uni-leipzig.de/index.php/mr2`, where data are made available together with details on the statistical analysis in the associated scholarly papers. With the data out in the open, readers will not only be able to evaluate for themselves the appropriateness of the analyses reported, but opportunities are created for improved analyses, either with current or with future statistical software.

## References

Arnon, I. and Ramscar, M. (2012). Granularity and the acquisition of grammatical gender: How order-of-acquisition affects what gets learned. *Cognition*, 122(3):292–305.

Baayen, R. H. (2010). A real experiment is a factorial experiment? *The Mental Lexicon*, 5(1):149–157.

Baayen, R. H. and Milin, P. (2010). Analyzing reaction times. *International Journal of Psychological Research*, 3:12–28.

Baayen, R. H., van Rij, J., de Cat, C., and Wood, S. N. (2015). Autocorrelated errors in experimental data in the language sciences: Some solutions offered by generalized additive mixed models. In Speelman, D., Heylen, K., and Geeraerts, D., editors, *Mixed Effects Regression Models in Linguistics*, page to appear. Springer, Berlin.

Balling, L. and Baayen, R. H. (2008). Morphological effects in auditory word recognition: Evidence from Danish. *Language and Cognitive Processes*, 23:1159–1190.

Balling, L. and Baayen, R. H. (2012). Probability and surprisal in auditory comprehension of morphologically complex words. *Cognition*, 125:80–106.

Balota, D., Cortese, M., Sergent-Marshall, S., Spieler, D., and Yap, M. (2004). Visual word recognition for single-syllable words. *Journal of Experimental Psychology: General*, 133:283–316.

Barr, D. J., Levy, R., Scheepers, C., and Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3):255–278.

Bates, D., Kliegl, R., Vasishth, S., and Baayen, R. H. (2015). Parsimonious mixed models.

Box, G. E. P. and Cox, D. R. (1964). An analysis of transformations (with discussion). *Journal of the Royal Statistical Society B*, 26:211–252.

Broadbent, D. (1971). *Decision and Stress*. Accademic Press, New York.

Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., and Munafò, M. R. (2013). Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5):365–376.

Cohen, J. (1983). The cost of dichotomization. *Applied Psychological Measurement*, 7:249–254.

De Vaan, L., Schreuder, R., and Baayen, R. H. (2007). Regular morphologically complex neologisms leave detectable traces in the mental lexicon. *The Mental Lexicon*, 2:1–23.

Dery, J. E. and Pearson, H. (2015). On experience-driven semantic judgments: A case study on the oneiric reference constraint. In *Proceedings of Quantitative Investigations in Theoretical Linguistics, Tübingen, 2015*.

Egly, R., Driver, J., and Rafal, R. D. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, 123:161–177.

Ernestus, M. and Cutler, A. (2015). BALDEY: A database of auditory lexical decisions. *The Quarterly Journal of Experimental Psychology*, (ahead-of-print):1–20.

Forster, K. (2000). The potential for experimenter bias effects in word recognition experiments. *Memory & Cognition*, 28:1109–1115.

Francis, G. (2012). Publication bias and the failure of replication in experimental psychology. *Psychonomic Bulletin & Review*, 19(6):975–991.

Francis, G. (2013). Publication bias in "Red, Rank, and Romance in Women Viewing Men" by Elliot et al. (2010). *Journal of Experimental Psychology: General*, 142(1):292–296.

Gałecki, A. and Burzykowski, T. (2013). *Linear mixed-effects models using R: A step-by-step approach*. Springer Science & Business Media.

Hastie, T. and Tibshirani, R. (1990). *Generalized Additive Models*. Chapman & Hall, London.

Kleinschmidt, D. F. and Jaeger, T. F. (2015). Robust speech perception: Recognize the familiar, generalize to the similar, and adapt to the novel. *Psychological Review*, 122(2):148.

Kliegl, R., Kuschela, J., and Laubrock, J. (2015). Object orientation and target size modulate the speed of visual attention. *Manuscript, University of Potsdam*.

Kliegl, R., Wei, P., Dambacher, M., Yan, M., and Zhou, X. (2011). Experimental effects and individual differences in Linear Mixed Models: Estimating the relationship between spatial, object, and attraction effects in visual attention. *Frontiers in Psychology*, 1:1–12.

Lin, X. and Zhang, D. (1999). Inference in generalized additive mixed models using smoothing splines. *JRSSB*, 61:381–400.

MacCallum, R., Zhang, S., Preacher, K., and Rucker, D. (2002). On the practice of dichotomization of quantitative variables. *Psychological Methods*, 7(1):19–40.

Marsolek, C. J. (2008). What antipriming reveals about priming. *Trends in Cognitive Science*, 12(5):176–181.

Masson, M. E. and Kliegl, R. (2013). Modulation of additive and interactive effects in lexical decision by trial history. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(3):898.

Matuschek, H., Kliegl, R., and Holschneider, M. (2015a). Smoothing spline ANOVA decomposition of arbitrary splines: An application to eye movements in reading. *PLoS ONE*, 10(3).

Matuschek, H., Kliegl, R., Vasishth, S., Baayen, R. H., and Bates, D. (2015b). Balancing Type I Error and Power in Linear Mixed Models.

Mulder, K., Dijkstra, T., Schreuder, R., and Baayen, R. H. (2014). Effects of primary and secondary morphological family size in monolingual and bilingual word processing. *Journal of Memory and Language*, 72:59–84.

Open Science Collaboration (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251):aac4716.

Pinheiro, J. C. and Bates, D. M. (2000). *Mixed-effects models in S and S-PLUS*. Statistics and Computing. Springer, New York.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32:3–25.

R Development Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Ramscar, M., Hendrix, P., Love, B., and Baayen, R. H. (2013). Learning is not decline: The mental lexicon as a window into cognition across the lifespan. *The Mental Lexicon*, 8:450–481.

Ramscar, M., Shaoul, C., and Baayen, R. H. (2015). Why many priming results don't (and won't) replicate: A quantitative analysis. Manuscript, University of Tübingen.

Ramscar, M., Yarlett, D., Dye, M., Denny, K., and Thorpe, K. (2010). The effects of feature-label-order and their implications for symbolic learning. *Cognitive Science*, 34(6):909–957.

Sander, P. and Sanders, L. (2006). Rogue males: Sex differences in psychology students. *Electronic Journal of Research in Educational Psychology*, 8(4):85–108.

Sanders, A. (1998). *Elements of Human Performance: Reaction Processes and Attention in Human Skill*. Lawrence Erlbaum, Mahwah, New Jersey.

Tabak, W., Schreuder, R., and Baayen, R. H. (2005). Lexical statistics and lexical processing: Semantic density, information complexity, sex, and irregularity in Dutch. In Kepser, S. and Reis, M., editors, *Linguistic Evidence — Empirical, Theoretical, and Computational Perspectives*, pages 529–555. Mouton de Gruyter, Berlin.

Tabak, W., Schreuder, R., and Baayen, R. H. (2010). Producing inflected verbs: A picture naming study. *The Mental Lexicon*, 5(1):22–46.

Taylor, T. E. and Lupker, S. J. (2001). Sequential effects in naming: A time-criterion account. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 27:117–138.

Ullman, M. T., Estabrooke, I. V., Steinhauer, K., Brovetto, C., Pancheva, R., Ozawa, K., and Mordecai, K. and Maki, P. (2002). Sex differences in the neurocognition of language. *Brain and Language*, 83:141–143.

van Rij, J., Baayen, R. H., Wieling, M., and van Rijn, H. (2015). itsadug: Interpreting time series, autocorrelated data using GAMMs. R package version 1.0.3.

Welford, A. (1980). Choice reaction time: Basic concepts. In Welford, A., editor, *Reaction Times*, pages 73–128. Accademic Press, New York.

Westfall, J., Kenny, D. A., and Judd, C. M. (2014). Statistical power and optimal design in experiments in which samples of participants respond to samples of stimuli. *Journal of Experimental Psychology: General*, 143(5):2020.

Wood, S. N. (2006). *Generalized Additive Models*. Chapman & Hall/CRC, New York.

Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society (B)*, 73:3–36.

Wood, S. N. (2013). On p-values for smooth components of an extended generalized additive model. *Biometrika*, 100:221–228.

Zhou, X., Chu, H., Li, X., and Zhan, Y. (2006). Center of mass attracts attention. *Neuroreport*, 17:85–88.

# Appendix

**The KKL dataset**

**random effects**

| Groups | Name | Std.Dev. | Corr |
|---|---|---|---|
| subj | Intercept | 0.161766 | |
| | Trial | 0.051180 | -0.273 |
| subj.1 | spt | 0.066809 | |
| subj.2 | grv | 0.033814 | |
| subj.3 | obj | 0.025609 | |
| subj.4 | orn | 0.076247 | |
| subj.5 | spt_orn | 0.033226 | |
| Residual | | 0.185591 | |

**fixed effects**

| | Estimate | Std. Error | $t$ value |
|---|---|---|---|
| Intercept | 5.659 | 0.018 | 322.91 |
| sze | 0.184 | 0.035 | 5.27 |
| spt | 0.074 | 0.008 | 9.62 |
| obj | 0.043 | 0.005 | 9.41 |
| grv | -0.001 | 0.005 | -0.17 |
| orn | 0.014 | 0.009 | 1.51 |
| Soa L | -0.010 | 0.001 | -12.56 |
| Soa Q | 0.019 | 0.001 | 20.61 |
| sze:spt | 0.048 | 0.015 | 3.14 |
| sze:obj | -0.012 | 0.009 | -1.30 |
| sze:grv | -0.037 | 0.010 | -3.66 |
| sze:orn | 0.039 | 0.018 | 2.20 |
| spt:orn | 0.020 | 0.006 | 3.12 |
| obj:orn | 0.009 | 0.007 | 1.26 |
| grv:orn | 0.011 | 0.007 | 1.52 |
| sze:spt:orn | -0.014 | 0.013 | -1.10 |
| sze:obj:orn | -0.003 | 0.014 | -0.24 |
| sze:grv:orn | -0.047 | 0.014 | -3.25 |
| Trial L | -0.043 | 0.006 | -7.47 |
| Trial Q | 0.015 | 0.001 | 16.96 |
| sze:Trial | 0.018 | 0.011 | 1.60 |
| orn:Trial | 0.028 | 0.003 | 8.64 |
| sze:TrialQ | -0.000 | 0.002 | -0.05 |
| orn:TrialQ | -0.006 | 0.005 | -1.27 |

Table 3: Summary of the LMM fitted to the KKL data set. Extensions to the reference model are highlighted. For factors, $-0.5/+05$ dummy coding was used. L: linear, Q: quadratic.

| A. parametric coefficients | Estimate | Std. Error | t-value | p-value |
|---|---|---|---|---|
| (Intercept) | 5.6818 | 0.0192 | 295.8684 | $< 0.0001$ |
| sze | 0.1832 | 0.0384 | 4.7703 | $< 0.0001$ |
| spt | 0.0744 | 0.0078 | 9.5469 | $< 0.0001$ |
| obj | 0.0418 | 0.0044 | 9.4394 | $< 0.0001$ |
| grv | -0.0005 | 0.0050 | -0.1030 | 0.9179 |
| orn | 0.0391 | 0.0157 | 2.4911 | 0.0127 |
| sze:spt | 0.0465 | 0.0156 | 2.9854 | 0.0028 |
| sze:obj | -0.0095 | 0.0089 | -1.0729 | 0.2833 |
| sze:grv | -0.0379 | 0.0100 | -3.8065 | 0.0001 |
| sze:orn | 0.0043 | 0.0314 | 0.1370 | 0.8910 |
| spt:orn | 0.0226 | 0.0065 | 3.4751 | 0.0005 |
| obj:orn | 0.0083 | 0.0070 | 1.1935 | 0.2327 |
| grv:orn | 0.0113 | 0.0070 | 1.6213 | 0.1050 |
| sze:spt:orn | -0.0148 | 0.0130 | -1.1420 | 0.2535 |
| sze:obj:orn | -0.0038 | 0.0139 | -0.2712 | 0.7862 |
| sze:grv:orn | -0.0471 | 0.0140 | -3.3658 | 0.0008 |
| B. smooth terms | edf | Ref.df | F-value | p-value |
| fs(Trial,subj) | 604.8040 | 772.0000 | 1241.7205 | $< 0.0001$ |
| re(subj,spt) | 76.9754 | 84.0000 | 28.4032 | $< 0.0001$ |
| re(subj,grv) | 46.1523 | 84.0000 | 1.9553 | $< 0.0001$ |
| re(subj,obj) | 35.8780 | 84.0000 | 1.8586 | $< 0.0001$ |
| re(subj,orn) | 54.7554 | 84.0000 | 1.9724 | $< 0.0001$ |
| re(subj,spt_orn) | 46.8990 | 84.0000 | 1.2980 | $< 0.0001$ |
| s(Trial): big+cardinal | 8.4515 | 8.6985 | 15.0442 | $< 0.0001$ |
| s(Trial): small+cardinal | 8.1975 | 8.5267 | 10.9101 | $< 0.0001$ |
| s(Trial): big+diagonal | 6.0122 | 6.6187 | 5.8248 | $< 0.0001$ |
| s(Trial): small+diagonal | 8.1585 | 8.5015 | 11.0372 | $< 0.0001$ |
| s(Soa) | 5.4968 | 6.6352 | 94.4320 | $< 0.0001$ |

Table 4: Summary of the full GAMM fitted to the `KKL` data set. Extensions to the reference model are highlighted. For factors, $-0.5/+05$ dummy coding was used. A separate factor was defined with four levels, one for each combination of `Size` and `Orientation`, and a thin plate regression spline (`s()`) was fitted for each of its four levels. (Thin plate regression splines are smoothers constructed from a weighted sum of simple regular parametric smooth functions with a penalty for wiggliness.) `re(X,Y)` denotes random intercepts in $Y$ for grouping factor $X$. The penalized factor smooth for subject (`fs(Trial, subj)` includes by-subject intercept calibration.

**The baldey dataset**

| | Estimate | Std. Error | t value | Pr(>\|t\|) |
|---|---|---|---|---|
| Intercept | -0.93 | 0.06 | -14.71 | < 0.0001 |
| gender=male | 0.20 | 0.09 | 2.24 | 0.025 |
| | edf | Ref.df | F | p-value |
| s(LemmaFreq):gender=female | 4.52 | 5.17 | 36.43 | < 0.0001 |
| s(LemmaFreq):gender=male | 2.60 | 2.94 | 45.88 | < 0.0001 |
| te(word_duration,Trial):gender=female | 7.18 | 8.09 | 26.07 | < 0.0001 |
| te(word_duration,Trial):gender=male | 6.66 | 7.77 | 15.06 | < 0.0001 |
| re(word) | 1676.26 | 2777.00 | 1.82 | < 0.0001 |
| re(word,gender) | 488.29 | 5543.00 | 0.11 | < 0.0001 |
| re(subject,word_duration) | 17.47 | 18.00 | 35.13 | < 0.0001 |
| fs(session,subject) | 152.55 | 178.00 | 46225.44 | < 0.0001 |

Table 5: Summary of the model fit to the inverse transformed auditory lexical decision latencies in the `baldey` megastudy ($\rho = 0.2$). Factors received treatment dummy coding, `s()` denotes a thin plate regression spline, and `te()` a tensor product smooth. (A tensor product smooth constructs a wiggly surface out of restricted cubic splines with a penalty for wiggliness.) `re(X)` denotes random intercepts for grouping factor $X$, and `re(X,Y)` specifies random slopes for $Y$ for grouping factor $X$. `fs()` denotes a penalized factor smooth.

**The poems dataset**

| | Estimate | Std. Error | $t$-value | Pr(>\|t\|) |
|---|---|---|---|---|
| Intercept | 6.01 | 0.02 | 334.65 | $< 0.0001$ |
| | edf | Ref.df | $F$ | $p$-value |
| te(Fre,Trial) | 10.43 | 11.56 | 79.20 | $< 0.0001$ |
| re(Poem) | 81.08 | 86.00 | 19.40 | $< 0.0001$ |
| fs(Trial, Subject) | 2269.74 | 2933.00 | 341.31 | $< 0.0001$ |
| re(Subject, Fre) | 304.18 | 325.00 | 14.98 | $< 0.0001$ |

Table 6: Summary of the generalized additive mixed model fitted to the `poems` data, with $\rho = 0.3$, and a tensor product smooth for `Frequency` by `Trial` (fREML 49300). `te(X,Y)` denotes a tensor product smooth, `re(X)` random intercepts for grouping factor $X$, `re(X,Y)` denotes random slopes for $Y$ by grouping factor $X$, and `fs()` denotes a penalized factor smooth.

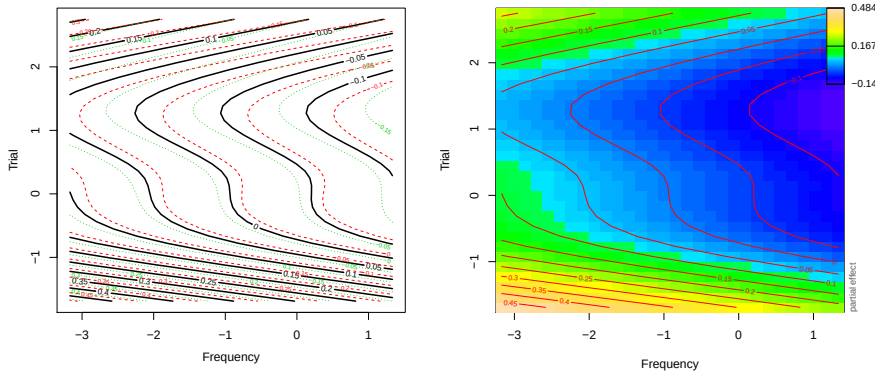


Figure 9: Tensor product smooth for the interaction of `Frequency` by `Trial` in the `poems` data set. In the left panel, green dotted lines indicate +1 standard error contour lines, and red dashed lines −1 standard error contour lines.

| | Estimate | Std. Error | $t$-value | Pr(>\|t\|) |
|---|---|---|---|---|
| Intercept | 6.04 | 0.02 | 334.97 | $< 0.0001$ |
| | edf | Ref.df | $F$ | $p$-value |
| ti(Fre) | 2.09 | 2.53 | 203.03 | $< 0.0001$ |
| ti(TrialSc) | 3.93 | 3.94 | 83.34 | $< 0.0001$ |
| ti(Fre,TrialSc) | 8.02 | 10.11 | 10.42 | $< 0.0001$ |
| re(Poem) | 81.08 | 86.00 | 19.36 | $< 0.0001$ |
| fs(TrialSc, Subject) | 2269.86 | 2933.00 | 337.14 | $< 0.0001$ |
| re(Subject, Fre) | 304.16 | 325.00 | 14.98 | $< 0.0001$ |

Table 7: Summary of a generalized additive mixed model fitted to the `poems` data, with additive effects of `Frequency`, `Trial`, and their interaction, and with $\rho = 0.3$ (fREML: 49307). `ti(X,Y)` denotes an independent tensor product smooth interaction term, and `ti(X)` the independent main effect of $X$. `re(X)` specifies random intercepts for grouping factor $X$, and `re(X,Y)` denotes random slopes for $Y$ by grouping factor $X$. `fs()` represents a penalized factor smooth, which absorbs the by-subject random intercepts.